Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-04 14:58:17
REFERENCE NO: 215

Author Names & Affiliations

- Ken Feldman - Univ. of Washington

Contact Email Address (for NSF use only)

(Hidden)

Research Domain, discipline, and sub-discipline

Marine Technician

Title of Submission

University National Oceanographic Laboratory System (UNOLS Satellite Network Advisory Group Response to NSF CI2030

Abstract (maximum ~200 words).

The University-National Oceanographic Laboratory System, (UNOLS) Satellite Network Advisory Group, (SatNAG) would like to provide this report to help guide the NSF in the area of cyber-infrastructure. The submitted document was written by Ken Feldman/Univ. of Washington, John Haverlack/Univ. of Alaska, Jin Meyer/Scripps Institution of Oceanography, and Laura Stolp/Woods Hole Oceanographic Institution. The submission has been completed by the UNOLS Office on 4 April 20917.

Question 1 Research Challenge(s) (maximum ~1200 words): Describe current or emerging science or engineering research challenge(s), providing context in terms of recent research activities and standing questions in the field.

Question 1
(maximum 1200 words) – Research Challenge(s). Describe current or emerging science or engineering research challenge(s), providing context in terms of recent research activities and standing questions in the field.

Productive, successful ocean science research increasingly requires fast, reliable, and robust connectivity to the Internet. The instrumentation and operating infrastructure have become more complex. The responsibilities of scientists are increasing, with greater demands on data, responsibilities, and outreach. These dependencies make functional Internet access mission critical.

The new research vessels (R/V Kilo Moana, R/V Sikuliaq, R/V Neil Armstrong, R/V Sally Ride) have been designed to carry fewer scientists. With fewer at-sea scientists, it is necessary that remote presence capabilities facilitate the ability for remote researchers to participate effectively in science operations and collaborate with others, be they on board or on shore 24 hours a day, 7 days a week. Real time, bidirectional high-bandwidth, and highly available telemetered data eases operational constraints in the field. Such data can also influence operational choices, in real time, from remote areas. Augmented bandwidth makes possible the ability to bring in experts from shore to diagnose issues, have shore support stand watches for data collection, and allow researchers to confer with other researchers (via Skype/Zoom/etc.) from the field to gain context for ongoing research.

NSF program requirements for outreach have also increased, and the ability to routinely bring the ship into the classroom in real time requires much more bandwidth than is currently possible. Currently, when Telepresence is included in a proposal, scientists must include extra costs for bandwidth upgrades. If the overall bandwidth for the fleet was increased to a functional baseline, the bandwidth can be reallocated instead of requiring a separate (expensive) contract on a per proposal basis.

The ability to transfer data to shore for calibration and artifact troubleshooting can be used to keep complex instruments running well. Current bandwidth constraints only allow snippets of data and the data can take days to reach shore, thus slowing the process of correction.

Multibeam data is a prime example. Science is requesting that the UNOLS vessels consistently run patch tests (yearly) to check the multibeam performance. The ability to move this data to shore for evaluation in a timely manner can resolve problems faster and more cost effectively.

The use of oceanographic observatories, including arrays of drones, drifters and other sensors in the field has increased. These instruments report into a central resource, via the Internet, to transfer data they are acquiring. Sometimes the only mechanism for command and control is via a satellite connection.

Multi-ship cruises, such as recent efforts with Armstrong/Endeavor/Sharp, and Ride/Sproul/Flip present operational challenges. The ability to share data between the ships is currently curtailed by the lack of bandwidth. Science is using what they have available, a multi-ship multi-cruise eventlog, but this has limited utility compared to sharing real time transmitted data about each ship's position, heading, and instrument information. Instantaneous collaboration amongst multi-ship experiments would be possible with more throughput.

The research vessels' science equipment is becoming more complex acquiring larger datasets. The ability to move more data to shore for quality control can save time and money, as will the ability for shore support to process data at sea in real-time to correct issues in the collection process.

The Internet of Things (IoT), defined as the interconnection via the Internet of computing devices embedded in everyday objects, enabling them to send and receive data, is causing issues of scale that our fleet has yet to acknowledge. A large increase in devices accessing the internet will constrain socket connection and Internet bandwidth resources.

Researchers and operators are presently in the impossible position of spending copious person-hours learning and/or teaching workarounds to cope with an underperforming set of satellite links compared to terrestrial broadband Internet. The scientific impact and utility of Research Vessels can be amplified many times over by enabling them to serve as observation hubs that serve data to a shoreside community of scientists working and teaching at their home institutions.

Question 2 Cyberinfrastructure Needed to Address the Research Challenge(s) (maximum ~1200 words): Describe any limitations or absence of existing cyberinfrastructure, and/or specific technical advancements in cyberinfrastructure (e.g. advanced computing, data infrastructure, software infrastructure, applications, networking, cybersecurity), that must be addressed to accomplish the identified research challenge(s).

Question 2
(maximum 1200 words) – Cyberinfrastructure Needed to Address the Research Challenge(s). Describe any limitations or absence of existing cyberinfrastructure, and/or specific technical advancements in cyberinfrastructure (e.g. advanced computing, data infrastructure, software infrastructure, applications, networking, cybersecurity), that must be addressed to accomplish the identified research challenge(s).

21st century vessels are increasingly requiring remote access over the Internet to monitor, maintain and control systems critical to the safe and functional operation of the vessel. To be competitive in the future, our ships need to be remotely available "telescopes" into the ocean that can be considered an almost-always available resource to the scientific community.

High-performance Internet is a requirement for modern research in almost all disciplines. Early career scientists have come of age in the Internet era. Their professional training is framed by data-intensive scientific practice. As such, almost all researchers now come to sea expecting to have all of their devices connected. Asking researchers to modify their Cyber-behavior and reduce their network footprint in the context of their research is a distraction from — as well as an obstacle to — the focus on their research. In short, by spending less money on operational costs today, we are really causing researchers to pay for that “savings” in their people time and research quality at sea. If we are to build successful infrastructure today that makes our research vessels innovative and flexible enough to support the next generation of research, reliable and performant bandwidth must be a priority from now on.

The UNOLS fleet needs, at minimum, a tenfold increase in bandwidth to enter the realm of functional performance. With such an increase in bandwidth, it will be possible to stand up a single ad-hoc telepresence event on the fly if something unusual occurs which will allow for an increasingly dynamic set of options for science parties to get the input and collaboration they need from shore in order to help maximize the objectives of the expedition.

In order to accommodate high performance satellite connectivity, we need to invest in infrastructure today. Our most performant link needs to have, or simulate, a clear sky view with no obstructions, so that this link can remain online and available with 99.9% confidence. Ocean and Global class ships need to operate in the 4-8GHz C-band range for maximum immunity to the effects of weather on radio signals. Ka (26-40GHz FleetExpress) should be on board for a functional backup, and L-band (1-2GHz -- FleetBroadband) as a means of contact for mission-critical data, in the most severe situations. Intermediate vessels should be handled on a case-by-case basis. Regional vessels should pursue a Ku (12-18GHz) and Ka blend for maximum reliability, at the scale of the vessel. We also have an Earth station and vessels that have been using the same gear for around a decade, which is the edge of a usable lifespan for such equipment. The time is near to maintain existing installations, bringing them inline with modern design requirements and functionality if we are interested in improving or maintaining their performance.

Modern IT systems and software increasingly assume that devices are always connected to broadband-speed Internet for maintenance such as License Validation, Cloud Storage, Security Updates, Server/Client Application Design Requirements. Not keeping systems patched and updated is problematic in today's ever increasing risk with regard to Cybersecurity concerns.

Computing resources aboard vessels that require continuous operation with little-to-no interruption need to be engineered as highly available systems with no single point of failure. This typically means a virtual machine infrastructure for mission critical services with a small Storage Area Network (SAN) and redundant disk controllers to keep the infrastructure stable. This has been industry practice in datacenters for years and it is time to consider parts of our research vessels in need of the stability of such Enterprise-grade equipment. Further, we are starting to see some applications that need in excess of 100 TB of storage over 90 days, which is firmly entering the realm where we face issues of data scale; it will not be tenable to consider such data volumes without a SAN along with enterprise storage solutions that can cope with the redundancy and input/output requirements of so much data. We should act soon to ensure reliable, resilient mass storage solutions are on board. There is also a need for shipboard Supercomputing resources to process and reduce large data sets into smaller near real time data products that can be sent to shore for analysis. As an alternative approach, with functional bandwidth it is possible for someone on shore to actually perform data processing, freeing the limited number of people and computing resources on board to focus on quality data acquisition. In short, cyberinfrastructure aboard research vessels needs to use modern well-defined, well-understood best practices from datacenters in order to carry out their mission.

Networking infrastructure on many vessels is old and outdated. Many ships have old multimode fiber with a back bone of <1Gbit/s and struggle to move HD video for science to access. Newer ships have 10-20 Gbit/s backbones that can move HD video around the ship with ease. In short, networks are becoming more complex and more science instruments are using Ethernet to control systems; ships’ infrastructure needs to be updated to accommodate. For example, installing a Video Matrix/Network for use by science to bring HD video from the ocean floor, to science on the vessel, to classrooms around the world would benefit a great deal of missions, but the investment needs to be made to bring ships up to that standard.

Current fleet connections make it difficult to perform remote IT maintenance and troubleshooting, often taking up to 10 times as long to complete a task from shore. More robust Internet connectivity would make remote IT access more viable, and provide much needed IT support on board without having to allocate a berth. Properly designed, most IT-level problems can be diagnosed and addressed from afar as long as there is sufficient network performance to do so. For the edge cases when hands need to be laid on equipment, personnel on board need to be able to interact with remote support staff in a brief but meaningful way to carry out effective repairs.

Cybersecurity is a strong concern. As the number of critical systems that are exposed to the Internet increases, it will be more and more important to proactively install software updates and maintain best cybersecurity practices. Any modern organization is faced with the challenge delivering functional IT services and protecting information resources and assets from malicious or unauthorized access. UNOLS vessels are no exception. Before assets can be protected against risks there must be a way of identifying and characterizing the probability of and impact of risk. Fortunately, the National Institute of Standards and Technology (NIST) publishes Federal Information Processing Standards Publications (FIPS) to provide guidance for Federal agencies pursuant to compliance with the Federal Information Security Management Act (FISMA). While UNOLS has not given specific guidance to NIST or FIPS, these documents provide a reasonable starting point for risk assessment. FIPS PUB 199 (https://doi.org/10.6028/NIST.FIPS.199)

Question 3 Other considerations (maximum ~1200 words, optional): Any other relevant aspects, such as organization, process, learning and workforce development, access, and sustainability, that need to be addressed; or any other issues that NSF should consider.

Question 3
(maximum 1200 words, optional) – Other considerations. Any other relevant aspects, such as organization, process, learning and workforce development, access, and sustainability, that need to be addressed; or any other issues that NSF should consider.

Scientists - especially early career scientists and their discipline-specific collaborators - are increasingly seeking to incorporate multiple data streams from various sources in near or real time. Their scientific practices are increasingly data intensive. The “system of systems” we seek to develop are to support these scientific users.

Information systems need regular maintenance and a life cycle maintenance plan. For most computing-based gear, a reasonable lifespan is 3-7 years. Buying enterprise grade cyberinfrastructure presents a number of advantages: it inherently implies maintenance, increases chances of replacement and/or repair worldwide, increases chances of off-ship expertise arriving to assist in foreign ports, plus has inherent migration paths for future and a life cycle maintenance plan. Thus, money spent up front on industry-standard computing systems saves time and effort in the long term.

Numerous and varied required information systems are rapidly becoming more complex aboard research vessels. Relying on off-ship expertise to diagnose problems is a rapidly increasing requirement. As both the ship and science systems become more complex, it is increasingly likely that personnel on board will not have the depth of knowledge to rapidly/adequately diagnose problems with such systems when a problem occurs. We already have an immediate and pressing need for remote information systems support and expertise for a number of disciplines; this need is very likely to increase as time goes on.

Necessary, remote ship operations support is a present and pressing reality. Among many other systems, even newer sewage treatment equipment being installed today must have remote diagnosis capability. Remote engine systems on ships delivered or refitted this decade need remote support/troubleshooting. The ability to diagnose engine issues while at sea increases safety and reduces potential expense of disrupting scientific operations by having to return to port to address minor problems.

More throughput means remote research support. Multi-PI cruises can grow from few to many by virtual attendance. Visibility into data collection in real-time by multiple interested parties assures better data quality.

Institutions and researchers alike are moving towards using cloud services for collaboration efforts; information in cloud services are a fact of life. Without an increase of throughput to research vessels at sea, cost-saving measures used by many institutions will cease to be functionally available in the field. Many institutions are also moving towards using online secure access for timecards, travel docs, etc. So basic ship operations will be critically compromised unless the bandwidth issue is also addressed for non-science bandwidth need.

Regarding bandwidth, based on the above, it should be clear that we are under-provisioned to keep both research and operations successful within UNOLS. That said, even with appropriately provisioned satellite communication paths, satellites do not have infinite availability; speeds per-ship will likely struggle to exceed something like a single cellular phone’s speed, while facing the demands of an office building. As such, even with an improved satellite communication path, we are likely to struggle with demand that outpaces supply for ships, due to carrier availability and physical limitations of radio communications over such distances. Due to this, committees such as SatNAG will remain useful in recommending management strategies to cope with this discrepancy. Put simply: tenfold bandwidth will not mean ships at sea can stop managing their link in some fashion, but it should significantly ease the degree to which management needs to happen such that researchers and operators are able to accomplish modern, basic Internet tasks with ease.

In the near future, technological solutions such as Low Earth Orbit (LEO) satellite constellations may offer game-changing levels of throughput (e.g. 500Mbps) with more compact technology that will ease network management workarounds even more. That said, technology like this is a few years away from being production-ready and even further from being well tested and a reliable performer aboard research vessels. While technology like this promises appealing performance, the need to make our research vessels more functional is immediate. Also, future technology that eclipses our current options does not negate the utility of our current, tested technology from being a viable and functional backup for the next 5-10 years. As any ship operator knows: backups are critical for the ambitious schedules UNOLS vessels face. It is therefore our recommendation that we pursue designing current technology to meet current needs, while keeping a watchful eye toward the future.

We need to collect and report state-of-health metrics on the networking and information systems we rely on for research. We are in the business of supporting scientific, quantifiable data; we should also be about the business of collecting quantifiable data about the state of health of our cyberinfrastructure. Collecting (and making available) metrics of this nature allows us to predict failures before they are emergencies and understand failure modes we did not predict after the fact. More stable infrastructure means more productive research.

IT Automation and Configuration Management must be used. For recurring problems, adopting standard solutions with automated configurations will save IT and operational staff much technical debt, as well as provide a more stable, predictable, and consistent experience. Efforts across the fleet to automate the configuration of common instrumentation would benefit overall UNOLS data quality.

As mentioned previously, on-ship expertise/presence will prove to be increasingly difficult to maintain. Due to the complexity of the problems unfolding, organization and direction is needed in order to keep functional all the aspects of cyberinfrastructure in need of attention/development. Therefore, a community resource such as a Center of Excellence (CoE) to aid seagoing research environments in the architecture, installation and maintenance plans of information systems seems appropriate to consider. With such a resource in place we can move forward and firmly root our research fleet in 21st century best practices and solutions.